Hindawi
Complexity
Volume 2022, Article ID 7109529, 14 pages
https://doi.org/10.1155/2022/7109529

Research Article

Ecological and Coevolutionary Dynamics in Modern Markets Yield Nonstationarity in Market Efficiencies

Colin M. Van Oort,[1,2,3] **John Henry Ring IV**,[2,3,4] **David Rushing Dewhurst**,[2,3,5] **Christopher M. Danforth**,[2,3] **and Brian F. Tivnan**[1,2,3]

[1]*The MITRE Corporation, McLean, Virginia 22102, USA*
[2]*Complex Systems Center, University of Vermont, Burlington, Vermont 05405, USA*
[3]*Computational Finance Lab, Burlington, Vermont 05405, USA*
[4]*MassMutual Data Science, Springfield, MA 01111, USA*
[5]*Charles River Analytics, Cambridge, MA 02138, USA*

Correspondence should be addressed to Colin M. Van Oort; cvanoort@mitre.org and Brian F. Tivnan; btivnan@mitre.org

Received 5 November 2021; Revised 19 December 2021; Accepted 22 December 2021; Published 2 June 2022

Academic Editor: Paulo Jorge Silveira Ferreira

The U.S. stock market is one of the largest and most complex marketplaces in the global financial system. Over the past several decades, this market has evolved at multiple structural and temporal scales. New exchanges became active, and others stopped trading, regulations have been introduced and adapted, and technological innovations have pushed the pace of trading activity to blistering speeds. These developments have supported the growth of a rich machine-trading ecology that leads to qualitative differences in trading behavior at human and machine time scales. We conduct a longitudinal analysis of comprehensive market data to quantify nonstationary dynamics throughout this system. We quantify the relationship between fluctuations in the number of active trading venues and realized opportunity costs experienced by market participants. We find that information asymmetries, in the form of quote dislocations, predict market-wide volatility indicators. Lastly, we uncover multiple micro-to-macro level pathways, including those exhibiting evidence of self-organized criticality.

1. Introduction

The U.S. National Market System (NMS) is composed of adaptive agents with goals and strategies for achieving them [1–3]. The NMS and its various participants continue to coevolve: trading strategies adapted, publicly traded companies are listed and delisted, regulations change, and stock exchanges launch and cease operations. This coevolution may yield emergent phenomena, such as bubbles and crashes [4–6].

The motivation for our study has two primary components. The theoretical component encompasses recent extensions to prevailing financial theory: from the Efficient Market Hypothesis to the Adaptive Market Hypothesis [1]. Market efficiency is tightly coupled to various aspects of the market ecology; namely, the number of competitors, the frequency and magnitude of profit opportunities, and the adaptability of market participants. The empirical component encompasses recent studies to characterize bubbles and crashes [6] and market inefficiencies [7, 8]. Central to both components, and core to our study, are nonstationary dynamics arising from information asymmetries as observed via a comprehensive order flow dataset.

The NMS has changed in many ways during our study period, which spans Q3 2009 through Q2 2017. Aside from the commonly observed fluctuations in asset price and volatility [9–12], we show that quote dislocations [7, 8] display nonstationary behavior that correlate with several market-wide volatility measures. Further up the hierarchy of market complexity, the topology of NMS infrastructure has changed. At the beginning of our study period, NMS infrastructure was primarily located in Carteret and Weehawken New Jersey. In 2010, NYSE relocated their

FIGURE 1: The U.S. national market system (NMS) is a complex, multilayered system exhibiting dynamic behaviors that vary based on network topology, active trading strategies, the regulatory environment, and more. The top panel shows several time series generated by the system, namely, the returns and volatility for the trading symbol PRU, along with an indicator for information asymmetry intensity. The second panel investigates higher-level system dynamics, displaying some of the network topologies that the NMS has assumed over the course of our study period (Q3 2009 – Q2 2017). The third panel shows a more detailed view of the NMS as it stood in September 2016, following IEX's transition from an alternative trading system to a national securities exchange. The final panel summarizes the activity profiles for all exchanges that functioned during our study period. Shaded dates indicate that an exchange was active, while unshaded dates indicate abnormal or unexpected inactivity. Expected outages, such as weekends and holidays, are not shown.

exchanges to a data center in Mahwah, and in 2016, IEX began operating an exchange in Weehawken. We show that changes in the topology of the NMS are associated with opportunity costs realized by NMS participants, even though the most rapidly fluctuating trading venues execute a comparatively small portion of all trades.

We briefly outline the structure of the NMS, which is summarized in Figure 1. The NMS is composed of many interacting components, including data centers, exchanges, traders, regulators, and multiple layers of regulations. The data centers that host the computing infrastructure of the NMS are located in north-eastern corner of New Jersey. Carteret houses the NASDAQ family of exchanges, Mahwah the NYSE family, and Secaucus the BATS and DirectEdge families, while Weehawken is home to the Investors Exchange (IEX). Traders that aim to minimize the latency of their connection to the NMS infrastructure often colocate their trading servers in the same data centers as the exchanges they often interact with. These data centers are connected by communication channels built with state-of-the-art technology, including fiber optic cables, laser free space optics, and millimeter wave wireless.

Market participants connected to the NMS can subscribe to a variety of data feeds offered by an exchange (direct feed), a Securities Information Processor (SIP feed), or a data vendor (third party feed). Direct feeds give traders access to top of book information (e.g., quote and trade messages), depth of book information (e.g., add, modify, and cancel messages), and administrative notifications (e.g., trading halt and order imbalance messages). SIPs are a market utility that aggregate information from exchanges, which introduces a small amount of latency relative to direct feeds but provide important market-wide signals. SIPs offer top of book information via a trade and quote (TAQ) feed, the National Best Bid and Offer (NBBO; an important market-wide best price signal), and Limit-Up/Limit-Down bands (LULD; a system to dampen excessive volatility). Trade-through protections mandate that trades execute at a price that is at least as good as the NBBO, giving it an importance beyond a simple price signal for traders.

Central to the structure of the NMS is the differentiation between National Securities Exchanges, or "lit" venues, and alternative trading systems (ATS), also known as "dark" venues, which are governed by different regulations and are subject to different reporting requirements [13]. National Securities Exchanges are required to provide quotes to the SIP, resulting in active participation in a market wide price discovery process. ATS are only required to issue quotes in extraordinary circumstances [13] and thus tend to impact price discovery indirectly.

The arrangement of data centers and exchanges has been fairly dynamic since NSEs began operating electronic trading systems. The transition to electronic trading was sparked by Nasdaq's acquisition of Inet in 2005 (originally located in Carteret) and was quickly followed by NYSEs acquisition Archipelago in 2006 (originally located in Weehawken). Since that time, the NMS has grown to include 12 NSEs and roughly 40 ATS distributed across four geographic locations. Births, deaths, technical glitches, and unplanned market events caused exchanges to switch on and off a total of 86 times during our study period, excluding normally scheduled outages (e.g. holidays).

We extend the analyses presented in Tivnan et al. [7] and Ring et al. [8] to cover more than 3000 trading symbols over a period of roughly 8 calendar years. We identify correlations between data feed information asymmetries, opportunity costs associated with those asymmetries, and trading venue fluctuations. While some of these correlations may be relatively easy to infer, they lead us to a less intuitive result. Metrics derived from data feed information asymmetries have forward predictive power with respect to several volatility measures. We also discuss long-term trends in each of the studied quantities and implications those trends may have.

2. Related Work

Our work aims to quantify and investigate aspects of market efficiency in the NMS. The prevailing theory is the Efficient-Market Hypothesis (EMH), which was popularized in its current form by [14] and supported by many including Bachelier [15], Mandelbrot [16], and Samuelson [17]. EMH states that asset prices reflect available information, and EMH is usually presented as having three forms. The three forms of EMH are differentiated by the information that each form claims is incorporated into asset prices. The strong form claims that asset prices incorporate all public and private information, the semistrong form only considers publicly available information, and the weak form only considers historical prices. Regardless of which form we inspect, the connection between the EMH, excess profits, and random walks follows from the definition. If the current price of an asset already accounts for a piece of information, then it should not have any bearing on future prices. If that information has no bearing on future prices, trading based on that information should not provide any excess returns.

Though the intuition behind the EMH is widely accepted, objections have arisen regarding its various formalizations. Several works have indicated that perfectly efficient markets are unlikely to exist, due to the cost of obtaining and acting upon relevant information, thus calling into question the realism of the strong form of EMH [18, 19]. The theoretical justification for the connection between EMH and random walk models has been contested, with some works supporting the connection [14, 17] and others refuting it [20, 21]. Likewise, a body of empirical work with conflicting findings exists. Empirical works in support of the EMH often focus on events where markets quickly react to information [22–26] or develop and apply statistical tests of efficiency [27–29]. Conversely, many studies have found evidence that tests for efficiency are intermittently passed at best [30–32] or that price anomalies arise consistently [33, 34].

Sustained controversy over the EMH suggests that market efficiency is not necessarily a static property and has led to the rise of alternative theories, such as the Adaptive Markets Hypothesis (AMH) [35]. With findings covering efficiency dynamics in emerging markets [36], established markets [37–39], and interasset variation [40], recent empirical work has largely supported the AMH. Rather than testing for the existence or absence of market efficiency, tools that quantify aspects of efficiency are needed to better understand the dynamics suggested by the AMH. Ding et al. [41] and Bartlett and McCrary [42] investigated quote

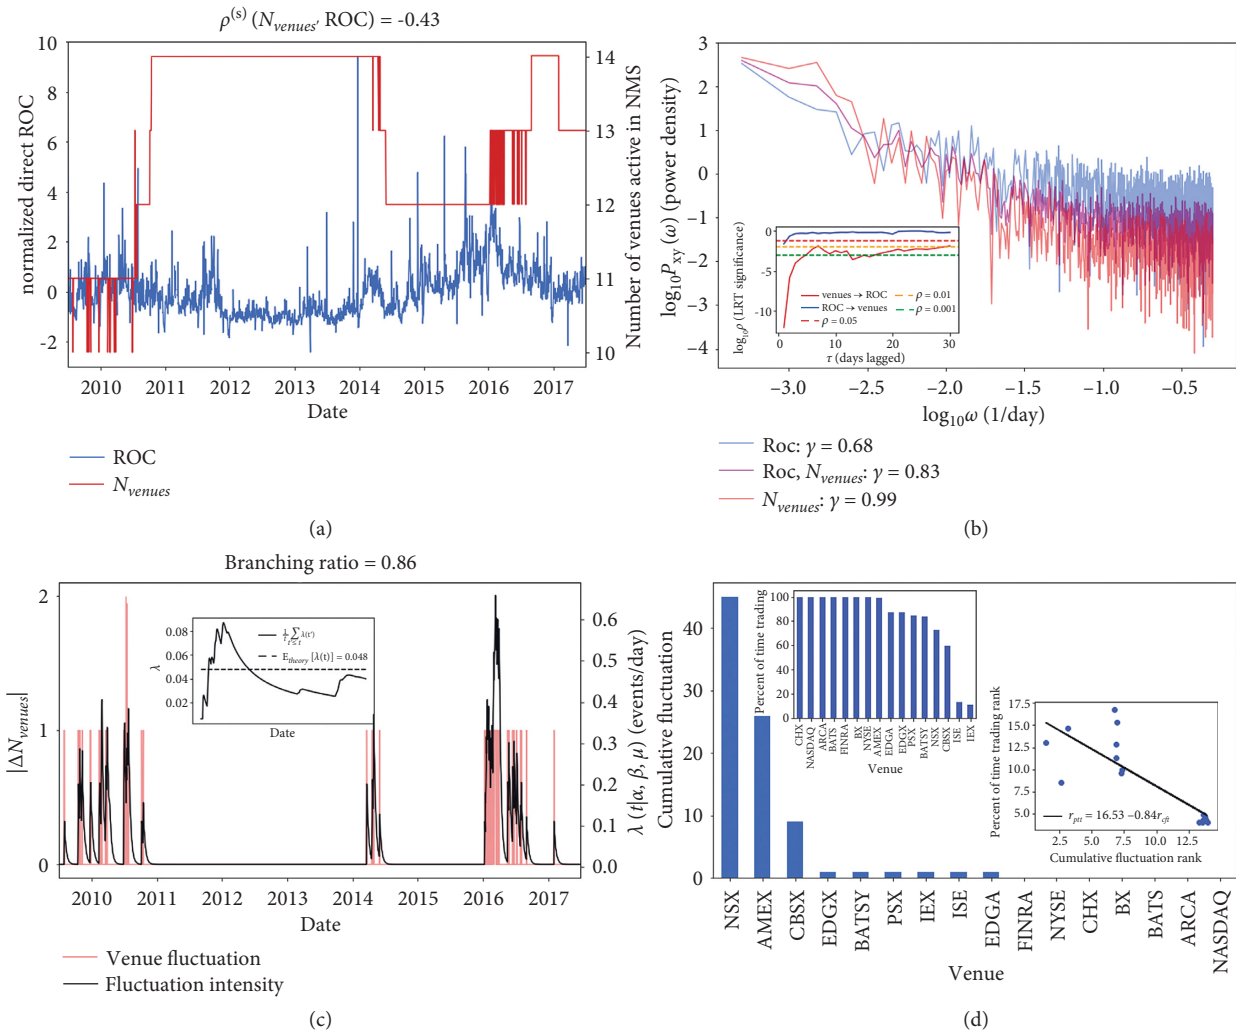

Figure 2: (a) The number of active trading venues and realized opportunity cost (ROC), normalized to have zero mean and unit variance, are negatively correlated. (b) Estimates of their power spectral densities allow for assessment of the stationarity of the time series. The inset demonstrates the granger-causal relationship between venues and ROC (venues $\longrightarrow$ ROC). (c) The time series of fluctuation in the number of trading venues and the latent intensity time series of a hypothesized Hawkes process model for venue fluctuation. The inset panel displays the convergence for this model. (d) Cumulative fluctuation for each exchange over the study period. The first inset plot shows the total percent of time under study that each venue was active, and the second inset plot displays a negative association between fluctuation and total percent of time that a venue was active.

dislocations between pairs of data feeds and provided insights into potentially pervasive information asymmetries. Wah [43] analyzed arbitrage opportunities, shining a light on the potential gross profits of certain trading strategies. Ahmed and Satchell [44] developed a bubble detection method that allows for effective interasset comparisons. Godfrey [45] studied a measure of efficiency based on the relative profits of ideal passive and active traders.

3. Methods

We analyzed data from Thesys Technologies [46] that contained every message from each SIP and direct feed in a unified data format. Consisting of ~40 PB of data, this dataset covered trading activity from ~2008 to present. Thesys collected this data using hardware colocated in the Carteret datacenter and applied a μ s-resolution timestamp to each message on receipt to mitigate clock synchronization issues. Thesys served as the data provider for the SEC's Market Information Data Analytics System (MIDAS) until 2019, when MayStreet acquired some of its assets and assumed its role as the MIDAS data provider.

Using Dislocation Segments (DSs) and Realized Opportunity Costs (ROC) associated with those DSs, as described in Tivnan et al. [7] and Ring et al. [8], we cataloged information asymmetries in the NMS and estimated their impacts. A DS is an information asymmetry between a pair of data feeds that is caused by quote price discrepancies. We follow the previous work in considering prices displayed by the SIP NBBO and a synthetic Direct Best Bid and Offer (DBBO). A DS begins when the price of the National Best Bid (NBB) or Offer (NBO) does not match its counterpart in

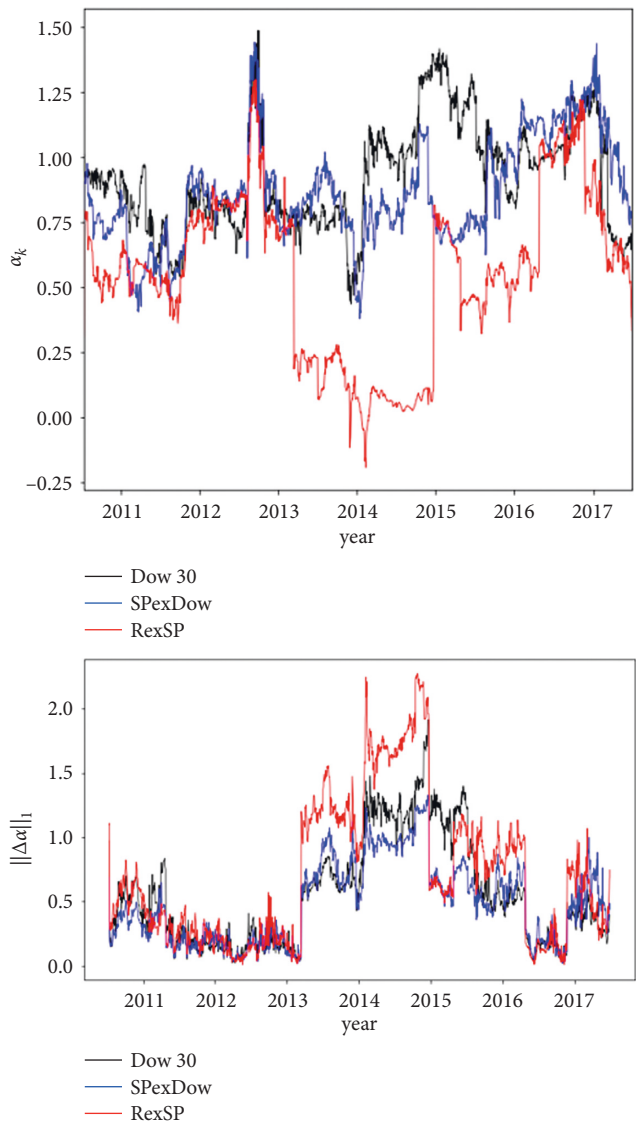

Figure 3: Spectral behavior of realized opportunity cost (ROC) time series is itself nonstationary. A 252-day (one market year) top-hat window τ_k was rolled over each time series of ROC corresponding with a mutually exclusive market segment and the within-window power spectral density was computed. These power spectra are universally well-fit by a function of the form $P_{xx}^{(\tau_k)}(\omega) \sim \omega^{-\alpha(\tau_k)}$.

the DBBO. An active DS ends if the prices converge, the relationship of the price divergence changes (i.e., NBO < DBO ⟶ DBO < NBO or vice versa), or the trading day ends. For each DS, we track the starting time, duration, maximum price divergence, and minimum price divergence. DSs are calculated independently for each trading symbol and each side of the market.

Trades that execute during DSs may incur opportunity costs, since routing and timing decisions may differ based on specifics of the available price signals. We quantify the Realized Opportunity Cost (ROC) for each trade as the number of shares traded × the total price difference between the NBBO and DBBO quotes on the appropriate side of the market. We only consider trades executed at one of the prices displayed by the prevailing NBBO. This provides a conservative estimate of total opportunity costs and focuses on trades that were likely informed by the NBBO.

Our analysis covers trading activity from June 1, 2009, to May 31, 2017, and all trading symbols that were included in the Russell 3000 at any point during that period. By covering both an extensive time period and a broad population of trading symbols, we aim to capture both longitudinal and cross-sectional relationships. We focus on the relationship between DSs and volatility of security prices as well as that between fluctuations in the number of exchanges and ROC. We primarily investigate direct ROC, opportunity costs that occur when a trade executes at a price displayed by the NBBO and that price was better for the active order than the corresponding price displayed by the DBBO. These are the costs that may be experienced by a participant that is an exclusive subscriber to direct feeds and does not consider SIP information when placing trades, or one who only considers lit liquidity.

Our methods build upon and extend those used by Ding et al. [41], Bartlett and McCrary [42], Tivnan et al. [7], Ring

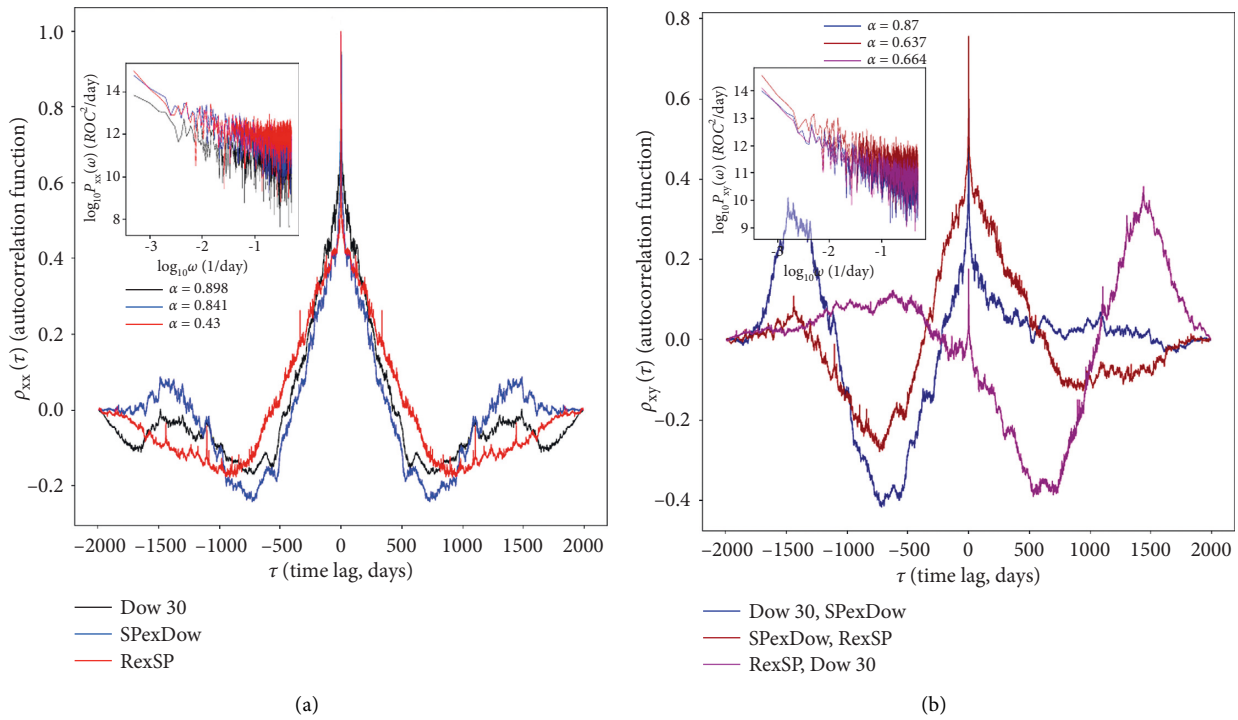

Figure 4: Autocorrelation (a) and cross-correlation (b) decay functions for realized opportunity cost (ROC) computed over three disjoint sets of trading symbols. ROC has a strong positive autocorrelation for hundreds of trading days, and two of the three sets of trading symbols exhibit similar cross-correlation behavior.

et al. [8], and to a lesser extent Wah [43]. Specifically, Ding et al. [41] studied similar quote divergences and used data collected from two distinct data feeds (i.e., SIP and direct), and that data was collected via a single, static observer. We advocate for this approach over that used by Bartlett and McCrary [42], which relied solely on SIP data, since this approach avoids potential issues with timestamp synchronization. Though our dislocation methodology and data quality are very similar to previous studies [7, 8, 41], the scope of our dataset is several orders of magnitude larger. Ding et al. [41] study dislocations in 24 securities across five exchanges and a single SIP for 16 days of trading. Tivnan et al. [7] expand this scope to consider dislocations and ROC in 30 securities across 13 exchanges and two SIPs for a year of trading. Ring et al. [8] go further, covering dislocations and ROC in more than 2900 securities for the same year. In this study, our analysis includes dislocations and ROC in more than 3000 securities across 14 exchanges and two SIPs spanning eight years of trading. Beyond extending the scope of these previous studies, we also apply new analytical tools, such as spectral methods and Granger causality, which we detail in the following section.

4. Results

Changes in NMS infrastructure had a material effect on observed direct ROC, which can be seen in Figure 2. The number of active exchanges ranged between a minimum of $N_{\text{venue}} = 10$ and a maximum of $N_{\text{venue}} = 14$ during the study period, with a total of $\sum_{t=1}^{T} |\Delta N_{\text{venues},t}| = 86$ fluctuations. The time series of venue fluctuations features long periods of stationarity interleaved with shorter periods of bursty activity.

We model this burstiness using an exponential Hawkes process of the form $|\Delta N_{\text{venues},t}| \sim \text{Poisson}(\lambda(t|\alpha,\beta,\mu))$, where $\lambda(t \mid \alpha,\beta,\mu) = \mu + \sum_{t_j < t} \alpha e^{-\beta(t-t_j)}$. When fit to the time series of venue fluctuations, this model has a branching ratio of $b = 0.86$, suggesting that if the underlying dynamics are accurately characterized by the model, 86% of venue fluctuations are endogenous and may occur as a result of another venue fluctuation. This endogeneity of bursty behavior, coupled with a pink noise power spectral density $S_{xx}^{(N_{\text{venues}})}(\omega) \sim \omega^{-0.99}$, suggests that the dynamics of trading venue fluctuations may be driven by self-organized criticality [47, 48]. There is an inverse correlation between a venue's total percent of within-study active days and the number of times that venue fluctuated on and off ($r = -0.82$, $p < 1 \times 10^{-4}$). This relationship is only partially explained by the birth and death of venues. A plurality of fluctuations are due to thinly traded venues switching on and off many times over relatively short periods without officially ceasing operations. ROC exhibits quasi-stationary spectral behavior throughout the study period, as seen in Figure 3, with power spectrum well-fit by $S_{xx}^{(\text{ROC})}(\omega) \sim \omega^{-0.68}$. However, the level values of ROC are negatively correlated with the number of active trading venues ($\rho^{(\text{Spearman})}(N_{\text{venues}}, \text{ROC}) = -0.43$). This indicates that the spectral characterization of ROC as simply colored noise does not capture the full extent of its dynamics. Granger-causality analyses conducted in both potential causality directions

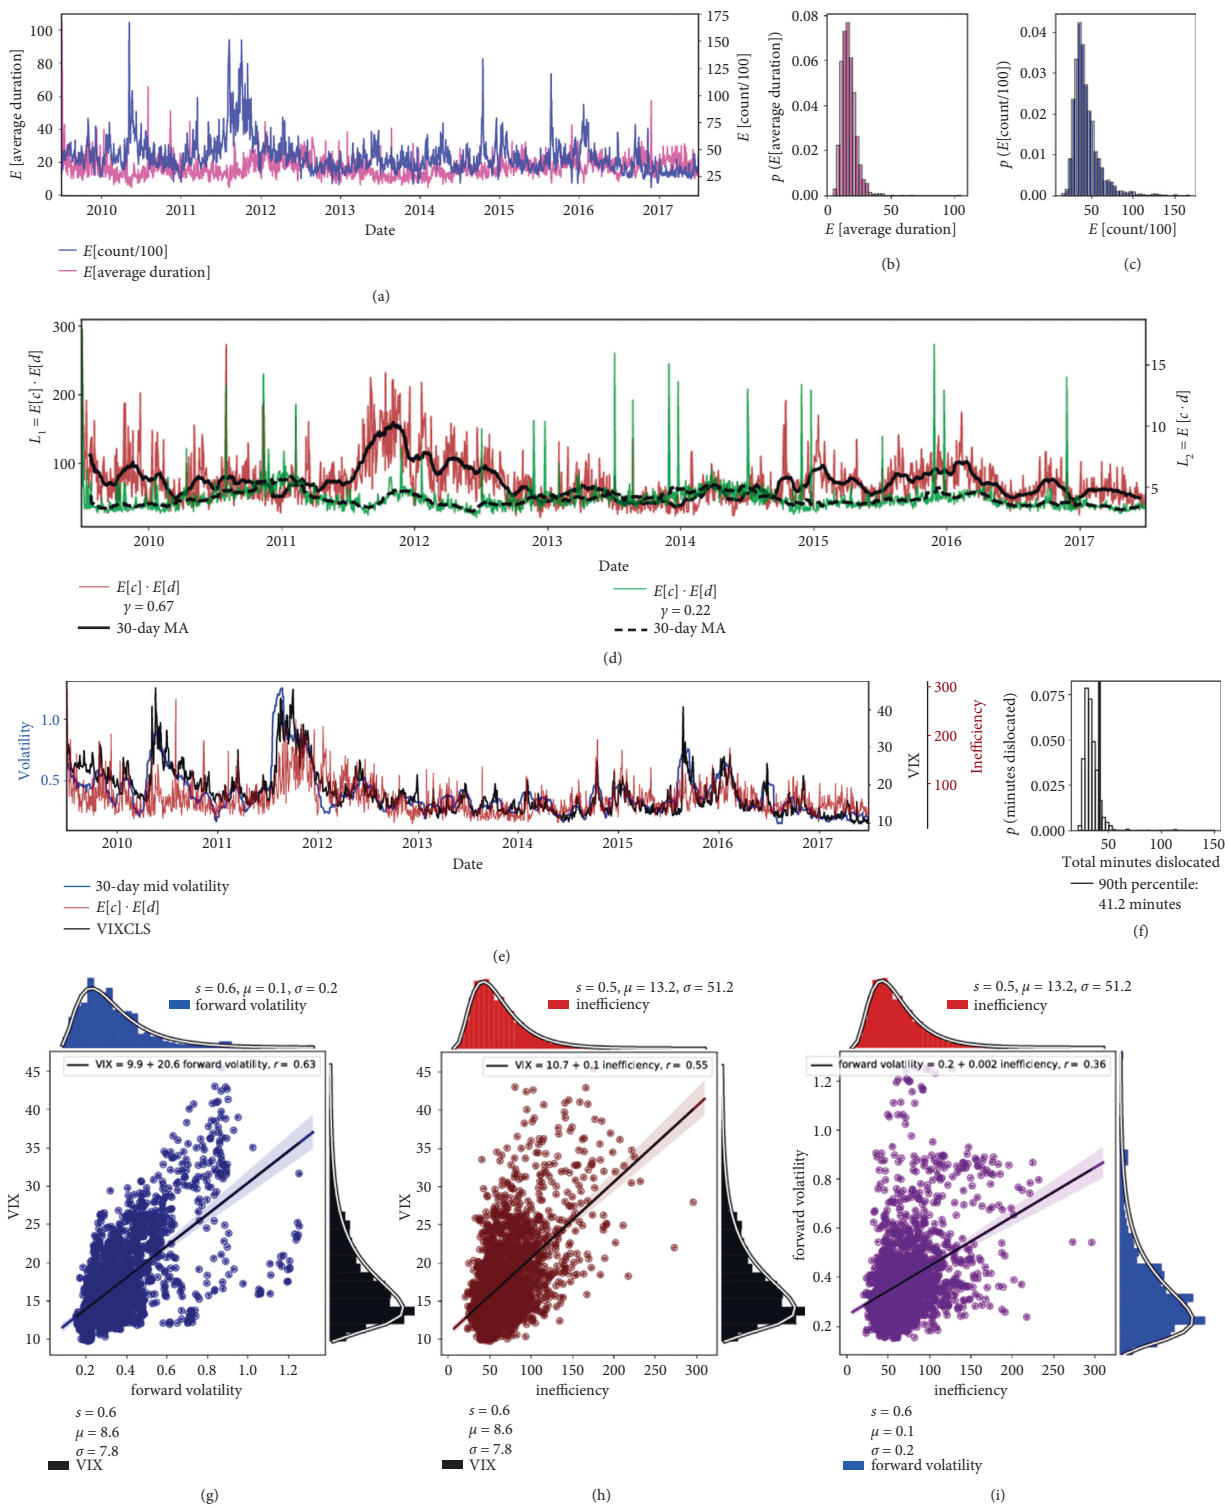

FIGURE 5: (a) Time series of average duration d and counts c of dislocation segments, which are computed and averaged across trading symbols for each day in the study period. (b, c) Histograms that capture the distribution of values assumed by the time series depicted in (a). (d) Two additional time series derived from the duration and count time series that provide an indicator for the intensity of information asymmetry that may have occurred on a particular trading day. These are defined as $\mathcal{I}_1 = E[c] \cdot E[d]$ and $\mathcal{I}_2 = E[c \cdot d]$. (e) The correlation between $\mathcal{I}_1$ and two measures of volatility. (f) The distribution of total dislocated minutes over trading symbols. (g–i) Joint distributions of forward volatility, VIX, and $\mathcal{I}_1$ along with marginal distributions of each statistic, which are well characterized by the lognormal probability density function.

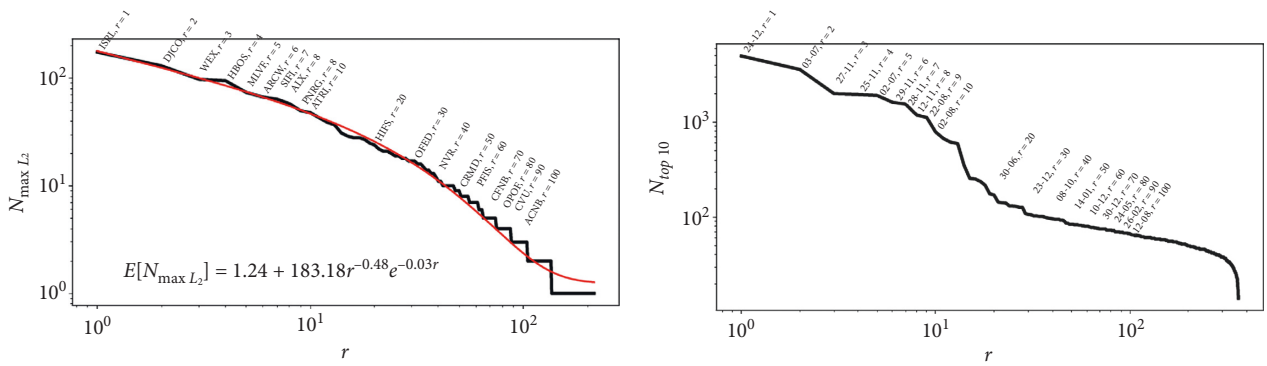

(a) (b)

Figure 6: Though $\mathcal{I}_2$ displays quasi-stationary spectral behavior, it also displays significant heterogeneity across trading symbols, as shown in (a) and regular seasonal spikes, as displayed in (b). For each trading symbol, we located the 10 days with the highest $\mathcal{I}_1$ values and incremented a counter associated with each, displayed as $N_{\text{top } 10}$ in (b) Of the top 10 of these days, seven are directly adjacent to a major holiday: Christmas (24–12), Thanksgiving (25–11, 27–11, 28–11, 29–11), and the anniversary of the U.S.'s founding (02–07, 03–07).

(venues $\longrightarrow$ ROC and ROC $\longrightarrow$ venues) highlight significant and persistent causality from venues to ROC (venues $\longrightarrow$ ROC). We show the relationship between the time lag τ and the significance of the test results $\log_{10} p$ in the upper left panel of Figure 2. By displaying the test significance for a range of τ we avoid arbitrary thresholds and unintended multiple comparisons. Additionally, Figure 4 shows that ROC occurring across disjoint groups of assets has nontrivial auto- and cross-correlation at long time lags, providing evidence of long memory in the generating process for ROC in addition to venue fluctuations.

Beyond their relationship with trading venue fluctuations, we are also interested in the longitudinal dynamics of dislocations, which can occur any time two or more trading venues and two information feeds are present [7]. In particular, we wish to understand how the prevalence of dislocations may have changed over time, along with how the duration and magnitude of dislocations may have developed over our period of study. Time series and time-decoupled statistics of these statistics are displayed in Figure 5. For each trading symbol x under study and each trading day T, we compute the total count $c_x(t)$ and average duration $d_x(t)$ of dislocations. From these, we calculate two indicators of expected dislocation intensity: $\mathcal{I}_1(t) = E_x[c_x(t)] \cdot E_x[d_x(t)]$ and $\mathcal{I}_2(t) = E_x[c_x(t) \cdot d_x(t)]$. Both indicators display quasi-stationary behavior over the study period ($S_{xx}^{(\mathcal{I}_1)}(\omega) \sim \omega^{-0.67}$, $S_{xx}^{(\mathcal{I}_2)}(\omega) \sim \omega^{-0.22}$). However, $\mathcal{I}_2$ displays seasonal behavior over the range of time studied, exhibiting large peaks on trading days adjacent to major U.S. federal holidays. Rank-intensity diagrams for $\mathcal{I}_1$ and $\mathcal{I}_2$, sorted by both date (modulo year) and trading symbol, are displayed in Figure 6.

$\mathcal{I}_1(t)$ is correlated with volatility statistics and has the ability to predict future volatility. We obtained three different volatility measures—the VIX volatility index, and 30-day midpoint and forward volatility of the Russell 3000 index (trading symbol: RUA). Midpoint volatility over τ time periods is defined as the standard deviation of log returns $r_{t'} = \log_{10}(X_{t'}/X_{t'-1})$ for $t' \in [t - \tau/2, t + \tau/2]$, while forward volatility is the same measure computed for $t' \in [t, t + \tau]$. For each volatility measure, we computed Pearson correlations between it and $\mathcal{I}_1$, finding significant correlation for all three combinations ($r(\mathcal{I}_1, \text{VIX}) = 0.55$, $r(\mathcal{I}_1, \text{mid vol}) = 0.47$, $r(\mathcal{I}_1, \text{forward vol}) = 0.36$). Since forward volatility is noncausal, correlation significantly greater than zero implies that $\mathcal{I}_1$ has some ability to predict forward volatility. Marginal distributions of volatility measures and our indicators are both well-fit by the lognormal family of distributions. Thus, we apply the SABR model [49], which assumes a time-dependent lognormal distribution whose volatility parameter varies following a geometric Brownian motion with zero-drift. The bottom three panels of Figure 5 highlight the correlation between $\mathcal{I}_1$ and volatility. This, coupled with lognormal marginal distributions, suggests that the pairs have similar generative mechanisms.

Quantile-based and distributional analysis of dislocation duration, shown in Figure 7, reveals complexity in the evolutionary dynamics of the NMS beyond what is seen in the time series of mean dislocation duration. The time series of median dislocation duration shows relatively smooth decay toward lower durations, punctuated by sharp drops that are likely due to upgrades in SIP technology [50, 51]. Distributions of dislocation duration quantiles have gradually drifted lower by nearly an order of magnitude. In contrast, distributions of mean dislocation duration have remained almost constant—expected due to their spectral characterization above—while distributions of dislocation duration standard deviations have generally increased over the study period. These observations indicate that the "typical" dislocation is becoming shorter, while the incidence of rare, long dislocations is growing more prevalent.

Emergent properties, such as realized opportunity cost and volatility, are driven by microlevel interactions among agents. These actions generate order flow, such as quotes and trades, which can then be used to predict macrolevel statistics. We have mainly focused on two micro-to-macro pathways, venue fluctuation to ROC and dislocation statistics to volatility. However, both dislocation and venue

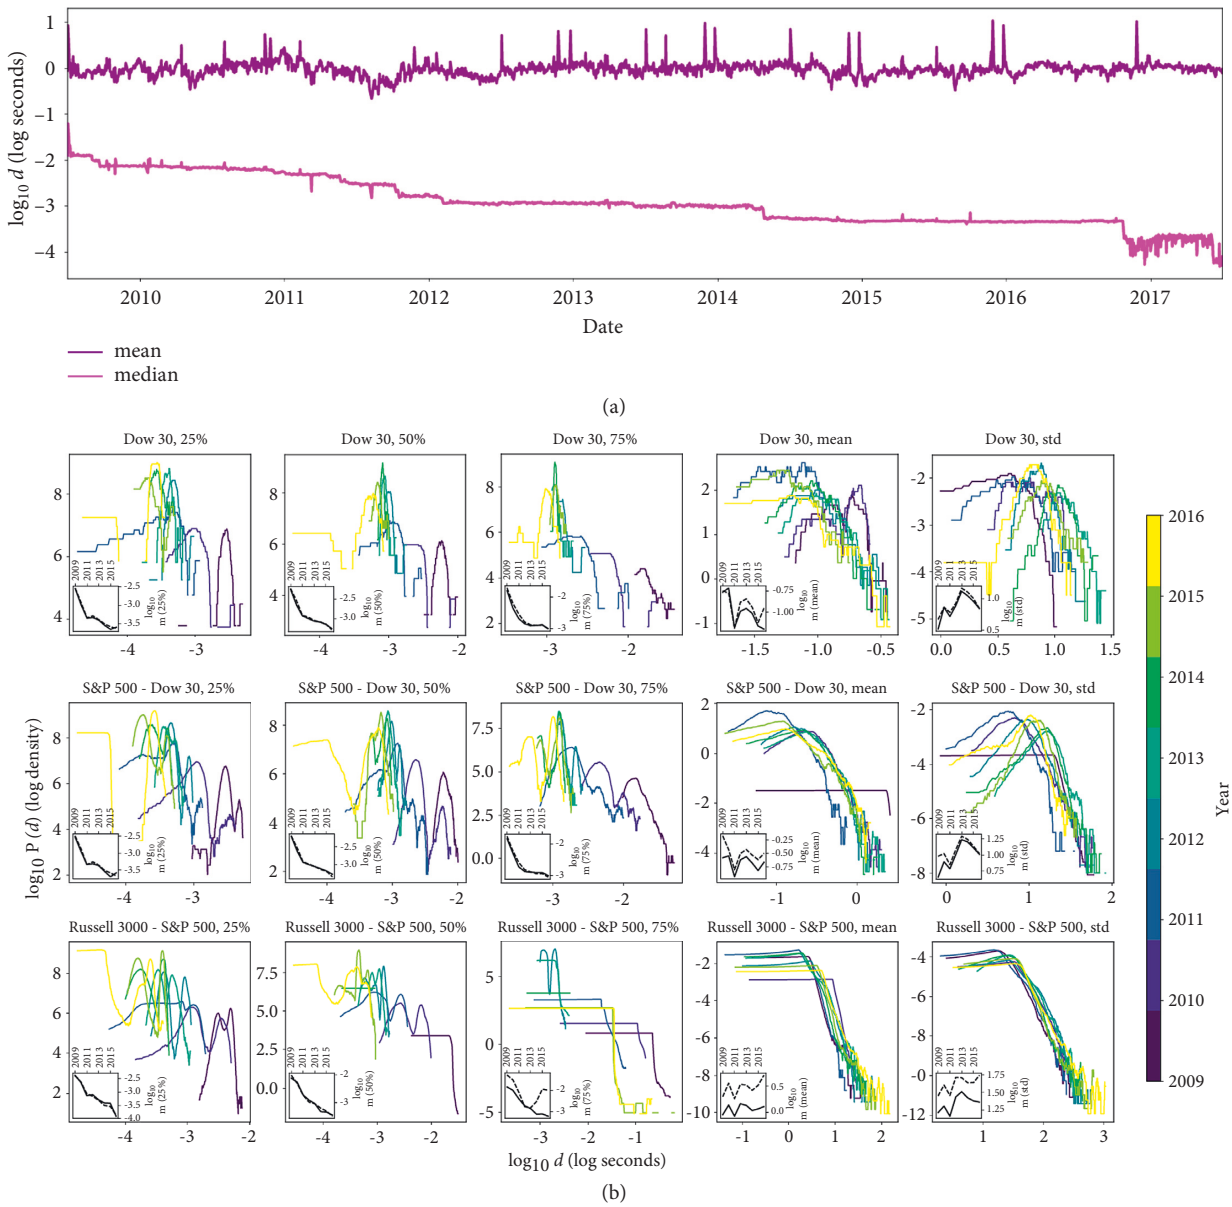

Figure 7: (a) Longitudinal dynamics in the mean and median dislocation duration. While the mean dislocation duration has remained relatively constant, the median has gradually decreased. The lower grid shows distributions of quantile and moment-based statistics over dislocation durations, for each calendar year and multiple disjoint sets of trading symbols. Distributions of quantile-based statistics (25th, 50th, and 75th percentiles) of dislocation duration declined by an order of magnitude over the study period, universally across the different sets of trading symbols, as displayed in the insets in (b). These two observations indicate that the "typical" dislocation duration has decreased, while the incidence of rare, long duration dislocations has increased.

Table 1: Linear regression diagnostics for the VIX as a function of trade statistics, dislocation statistics, and venue fluctuation.

Dep. variable	VIX		R-squared		0.738	
Model	OLS		Adj. R-squared		0.738	
F-statistic	1121.0		Prob (F-statistic)		0.00	
	Coef.	Std. error	t	$p > \|t\|$	[0.025]	[0.975]
Intercept	−0.0007	0.011	−0.064	0.949	−0.023	0.022
Differing trades	1.3024	0.058	22.378	<0.001	1.188	1.417
Count	0.6921	0.014	51.221	<0.001	0.666	0.719
Duration	0.0551	0.012	4.616	<0.001	0.032	0.079
Differing traded value	−1.4980	0.056	−26.834	<0.001	−1.607	−1.389
Venue	−0.1900	0.012	−15.532	<0.001	−0.214	−0.166

Table 2: Linear regression diagnostics for midvolatility as a function of trade statistics, dislocation statistics, and venue fluctuation.

Dep. variable	Midvolatility		R-squared		0.561	
Model	OLS		Adj. R-squared		0.560	
F-statistic	507.2		Prob (F-statistic)		0.00	
	Coef.	Std. error	t	$p > \|t\|$	[0.025]	[0.975]
Intercept	−0.0005	0.015	−0.031	0.976	−0.030	0.029
Differing trades	1.0034	0.075	13.304	<0.001	0.856	1.151
Count	0.638	0.018	36.198	<0.001	0.599	0.668
Duration	0.0254	0.015	1.638	0.102	−0.005	0.056
Differing traded value	−1.0967	0.072	−15.161	<0.001	−1.239	−0.955
Venue	−0.1172	0.016	−7.395	<0.001	−0.148	−0.086

Table 3: Linear regression diagnostics for forward volatility as a function of trade statistics, dislocation statistics, and venue fluctuation.

Dep. variable	Forward volatility		R-squared		0.392	
Model	OLS		Adj. R-squared		0.390	
F-statistic	255.9		Prob (F-statistic)		$2.10e-211$	
	Coef.	Std. error	t	$p > \|t\|$	[0.025]	[0.975]
Intercept	0.0002	0.018	0.011	0.991	−0.034	0.035
Differing trades	0.5046	0.089	5.686	<0.001	0.331	0.679
Count	0.5710	0.021	27.711	<0.001	0.531	0.611
Duration	0.0035	0.018	0.192	0.847	−0.032	0.039
Differing traded value	−0.6402	0.085	−7.520	<0.001	−0.807	−0.473
Venue	−0.1498	0.019	−8.032	<0.001	−0.186	−0.113

fluctuation statistics have predictive power with respect to ROC and volatility. We fit ordinary least squares models of the form

$$E[Y] = \beta_1 \text{differing trades} + \beta_2 \text{count} + \beta_3 \text{duration} + \beta_4 \text{differing traded value} + \beta_5 \text{venue}, \quad (1)$$

where Y is one of ROC, VIX, mid-volatility, or forward volatility. All variables were normalized as $z \longmapsto z - E[z]/\text{Std}(z)$. The design variables differing trades and differing traded value are control variables that measure, respectively, the number of trades that occurred during dislocations and the total value (share price × number of shares) exchanged by differing trades. Results of these regressions are summarized in Tables 1–4. The fraction of explained variance is high for both ROC ($R^2_{\text{adj}} = 0.795$) and VIX ($R^2_{\text{adj}} = 0.738$) and moderate for both mid ($R^2_{\text{adj}} = 0.560$) and forward ($R^2_{\text{adj}} = 0.390$) volatility. As detailed above, a greater number of trading venues is associated with lower ROC and volatility statistics in all models ($\partial Y/\partial \text{venue} < 0$, $p < 0.001$ for all Y), while a greater number of dislocations is associated with higher ROC and volatility statistics in all models ($\partial Y/\partial \text{count} > 0$, $p < 0.001$ for all Y). Average dislocation duration does have a significant ($p < 0.001$) positive effect on ROC ($\partial\text{ROC}/\partial\text{duration} = 0.1199$) and VIX ($\partial\text{VIX}/\partial\text{duration} = 0.0551$) but not on mid ($\partial\text{mid volatility}/\partial\text{duration} = 0.0254$, $p = 0.102$) and forward ($\partial\text{forward volatility}/\partial\text{duration} = 0.0035$, $p = 0.847$) volatility. These regression analyses are summarized by Figures 8 and 9.

5. Discussion

Our results indicate coherent nonstationarity in the U.S. National Market system. The topology of the system exhibits long periods of static behavior interspersed with short windows of rapid fluctuation in the number of active trading venues. This fluctuation, though infrequent, has a pronounced effect on dislocations and realized opportunity costs. Namely, realized opportunity costs decrease as the number of venues increases, while an increase in the number of venues Granger-causes future realized opportunity cost.

In addition to topological considerations, we examine the properties of dislocation segments longitudinally. Quantile- and moment-based analyses of distributions of dislocation segment durations uncover complimentary narratives on the evolution of the NMS. Distributions of 25th, 50th, and 75th percentiles of dislocation length show a near-monotone decrease toward lower values, indicating that market efficiency has generally improved over time. On the other hand, distributions of mean and standard deviation of dislocation duration remain nearly constant over the study period. In conjunction with decreasing quantile statistics, this implies an increase in the probability of large dislocation events.

We investigate indicators composed of aggregate dislocation statistics, which are positively correlated with multiple measures of current and future volatility. These indicators also highlight seasonal effects, including well documented irregularities that occur near U.S. holidays [52–55]. Linking the venue, ROC and dislocation, volatility

Table 4: Linear regression diagnostics for realized opportunity cost as a function of trade statistics, dislocation statistics, and venue fluctuation.

Dep. variable	ROC		R-squared		0.796	
Model	OLS		Adj. R-squared		0.795	
F-statistic	1549.0		Prob (F-statistic)		0.00	
	Coef.	Std. error	t	$p > \|t\|$	[0.025]	[0.975]
Intercept	$-1.138e-15$	0.010	$-1.12e-13$	1.000	−0.020	0.020
Differing trades	0.6078	0.051	11.836	<0.001	0.507	0.709
Count	0.0994	0.012	8.341	<0.001	0.076	0.123
Duration	0.1199	0.011	11.380	<0.001	0.099	0.141
Differing traded value	0.2093	0.049	4.249	<0.001	0.113	0.306
Venue	−0.1046	0.011	−9.692	<0.001	−0.126	−0.083

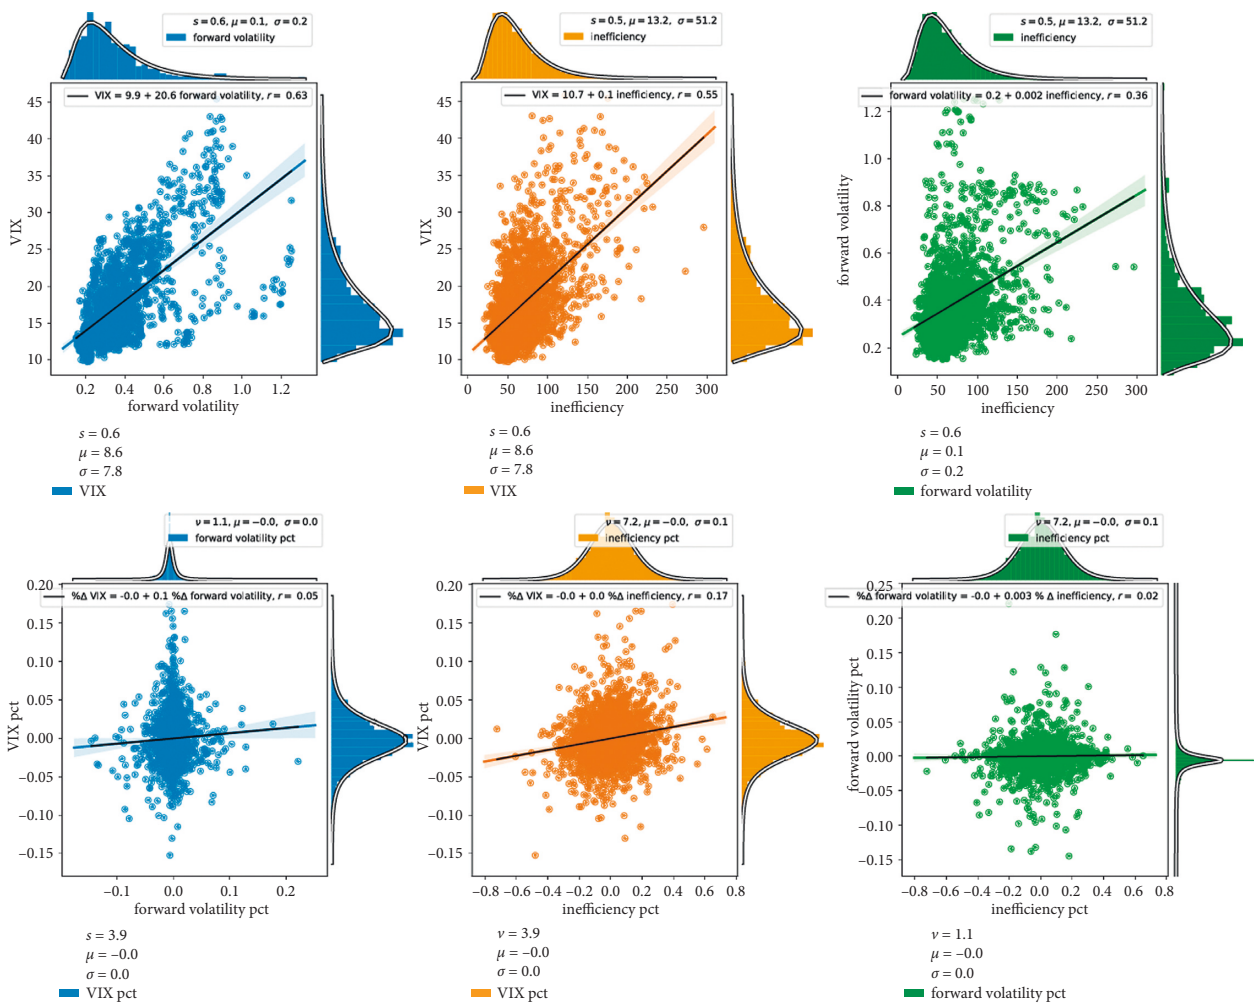

Figure 8: Common statistical measures of security price volatility, such as the VIX or rolling standard deviation of returns, are correlated with $\mathcal{I}_1$.

analyses, we find that venue fluctuation and dislocation statistics predict both ROC and price volatility of the Russell 3000 in least-squares regression models.

Though our work extends an existing line of inquiry to an unprecedented scale, it is not without limitations. Dislocation Segments, the primary unit under study, capture a specific form of information asymmetry and thus only capture a portion of the greater notion of market efficiency. Realized Opportunity Costs, which are derived from DSs and estimate their economic impact, might not provide an exhaustive estimate of economic impact. A more precise measure would likely account for available liquidity during each DS and a probabilistic estimate of whether a trader could access that liquidity. Additionally, neither DSs nor ROC are numerically well bounded, which may complicate comparisons between data from different assets or time

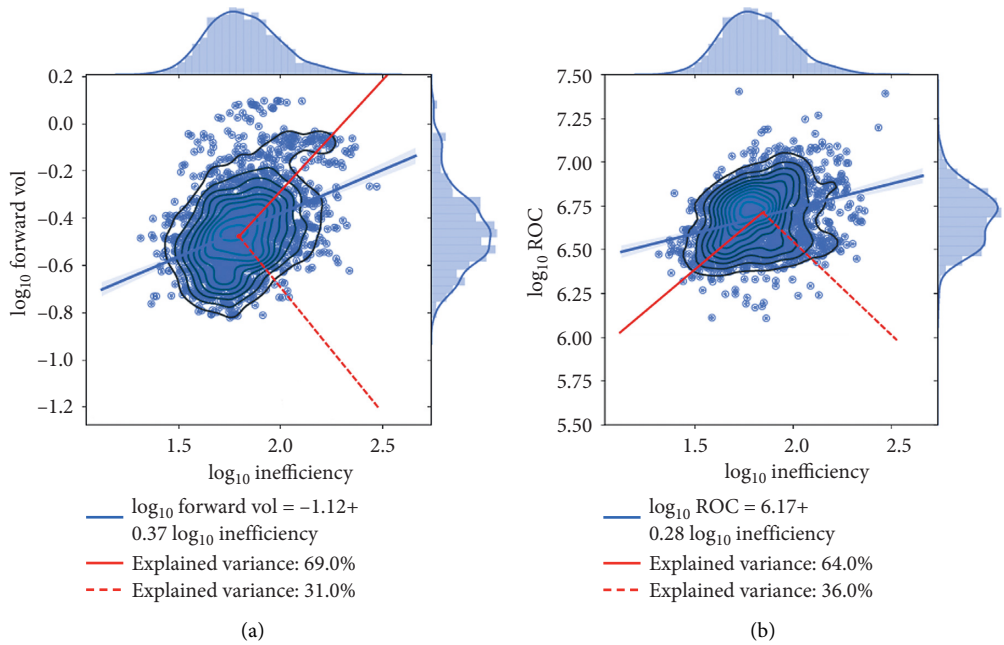

Figure 9: Regression results for log $\mathcal{I}_1$ vs. log forward volatility (a) and log $\mathcal{I}_1$ vs. log realized opportunity cost (ROC) (b). The solid red line shows the first principal component, the dashed red line shows the second principal component, and the black lines are density contours of the data. The fit line is shown in blue, and its equation is shown in the legend. These panels highlight the strong correlation between $\mathcal{I}_1$, forward volatility, and ROC.

periods. Finally, though the scope of our data allows us to investigate an impressive amount of cross-sectional variation and longitudinal dynamics, the level of aggregation needed to summarize and display the results may mask additional structures of interest.

6. Conclusion

We summarize the three salient aspects of our contributions to measure nonstationarity in the US National Market System (NMS). First, we identify predictive relationships between various market measures, to include dislocations predicting future volatility. Second, our findings provide evidence of increasing market efficiency via decreases in dislocation durations, coupled with increased tail risk of anomalous dislocations. Third, we uncover multiple micro-to-macro pathways, to include those exhibiting evidence of self-organized criticality.

Our methods, and their limitations, unearth several fruitful avenues for future work. There are opportunities to improve or extend the calculation and application of Dislocation Segments and Realized Opportunity Costs, as noted in our discussion of limitations. In contrast to our analysis, which investigated coarse-grained and large-scale dynamics, it may be informative to investigate the behavior of Dislocation Segments and Realized Opportunity Costs in close proximity to market anomalies such as the Flash Crash or the more recent activity in Game Stop. While we noted several alternative methods for quantifying market (in) efficiencies in our discussion of related work, this list of alternatives is not exhaustive. Therefore, a comparative study of various quantitative (in) efficiency measures would give researchers and policy makers a deeper understanding of their relative merits. Additionally, existing quantitative (in) efficiency measures each have limitations that present opportunities for new measures that fill different niches or to subsume existing measures.

While our findings highlight nonstationarity in the NMS throughout our period of study, the NMS continues to evolve. Notably, three new National Securities Exchanges became active in 2020 (Members Exchange (MEMX), Long-Term Stock Exchange (LTSE), and MIAX Pearl Equities (MIAX)). To better understand the dynamics of the current NMS, and how coevolutionary adaptations may impact those dynamics, we advocate for additional investigations of micro-behavior and micro-to-macro mechanisms. Principled models that consider market microstructure theory [56–58] and account for agent heterogeneity seem promising. The quantities observed in this study, dislocations and realized opportunity costs, along with their noted relationships with various econometrics, can serve as empirical targets for the next generation of market models.

Data Availability

Our data was originally provided by Thesys Technologies (https://www.thesystech.com/). Some components of Thesys have since been acquired by MayStreet (https://maystreet.com/). These data are publicly available as a commercial product.

Disclosure

All opinions and remaining errors are the sole responsibility of the authors and do not reflect the opinions nor perspectives of their affiliated institutions nor those of any funding agencies.

Conflicts of Interest

The authors declare that there are no conflicts of interest regarding the publication of this article.

Acknowledgments

The authors are grateful for discussions with Anshul Anand, James Bagrow, David Bringle, Eric Budish, Carl Burke, Peter Carrigan, Bill Gibson, Matthew Koehler, Blake LeBaron, Matthew McMahon, Mark Phillips, Mark Rosenthal, Wade Shen, David Slater, Jonathan Smith, and Jason Veneman. C.V.O., J.H.R., D.R.D., and B.F.T. were supported by the Defense Advanced Research Projects Agency (DARPA). C.M.D. was supported in part by a gift from the Massachusetts Mutual Life Insurance Company.

References

[1] A. W. Lo, *The Adaptive Markets Hypothesis*, Princeton University Press, Princeton, NJ, USA, 2019.

[2] M. Butler and D. Kazakov, "Testing implications of the adaptive market hypothesis via computational intelligence," in *Proceedings of the 2012 IEEE Conference OnComputational Intelligence For Financial Engineering & Economics (CIFEr)*, pp. 1–8, IEEE, New York, NY, USA, March, 2012.

[3] A. Urquhart and F. McGroarty, "Calendar effects, market conditions and the adaptive market hypothesis: evidence from long-run us data," *International Review of Financial Analysis*, vol. 35, pp. 154–166, 2014.

[4] J. D. Farmer and A. W. Lo, "Frontiers of finance: Evolution and efficient markets," *Proc Natl Acad Sci U S A*, vol. 96, no. 18, pp. 9991-9992, 1999.

[5] N. Goldenfeld and L. P. Kadanoff, "Simple lessons from complexity," *Science*, vol. 284, no. 5411, pp. 87–89, 1999.

[6] N. Johnson, G. Zhao, E. Hunsader et al., "Abrupt rise of new machine ecology beyond human response time," *Scientific Reports*, vol. 3, no. 1, pp. 1–7, 2013.

[7] B. F. Tivnan, D. R. Dewhurst, C. M. Van Oort et al., "Fragmentation and inefficiencies in us equity markets: evidence from the dow 30," *PLoS One*, vol. 15, no. 1, pp. 1–24, 2020.

[8] J. H. Ring, C. M. Van Oort, D. R. Dewhurst, T. J. Gray, C. M. Danforth, and B. F. Tivnan, "Scaling of Inefficiencies in the Us Equity Markets: Evidence from Three Market Indices and More than 2900 Securities," 2019, https://arxiv.org/abs/1902.04691.

[9] C. Brooks and M. J. Hinich, "Episodic nonstationarity in exchange rates," *Applied Economics Letters*, vol. 5, no. 11, pp. 719–722, 2010.

[10] R. L. Costa and G. L. Vasconcelos, "Long-range correlations and nonstationarity in the brazilian stock market," *Physica A: Statistical Mechanics and its Applications*, vol. 329, no. 1-2, pp. 231–248, 2003.

[11] T. Mikosch and C. Stărică, "Nonstationarities in financial time series, the long-range dependence, and the igarch effects," *Review of Economics and Statistics*, vol. 86, no. 1, pp. 378–390, 2004.

[12] T. A. Schmitt, D. Chetalova, R. Schäfer, and T. Guhr, "Nonstationarity in financial time series: generic features and tail behavior," *EPL (Europhysics Letters)*, vol. 103, no. 5, Article ID 58003, 2013.

[13] Securities and Exchange Commission, "Regulation of Nms Stock Alternative Trading Systems," 2018, https://www.sec.gov/rules/final/2018/34-83663.pdf, B-330281.

[14] E. F. Fama, "Efficient captial markets: a review of theory and emprical work," *The Journal of Finance*, vol. 25, no. 2, pp. 383–417, 1970.

[15] L. Bachelier, "Théorie de la spéculation," *Annales scientifiques de l'École normale supérieure*, vol. 17, pp. 21–86, 1900.

[16] B. B. Mandelbrot, "The variation of certain speculative prices," in *Fractals And Scaling in Finance*, pp. 371–418, Springer, New york, NY, USA, 1997.

[17] P. A. Samuelson, "Proof that properly anticipated prices fluctuate randomly," in *The World Scientific Handbook of Futures Markets*, pp. 25–38, World Scientific, Singapore, 1973.

[18] S. J. Grossman and J. E. Stiglitz, "On the impossibility of informationally efficient markets," *The American Economic Review*, vol. 70, no. 3, pp. 393–408, 1980.

[19] M. O. Jackson, "Efficiency and information aggregation in auctions with costly information," *Review of Economic Design*, vol. 8, no. 2, pp. 121–141, 2002.

[20] S. F. LeRoy, "Risk Aversion and the Martingale Property of Stock Prices," *International Economic Review*, vol. 14, pp. 436–446, 1973.

[21] R. E. Lucas Jr., "Asset Prices in an Exchange Economy," *Econometrica*, vol. 46, pp. 1429–1445, 1978.

[22] J. A. Busse and T. C. Green, "Market efficiency in real time," *Journal of Financial Economics*, vol. 65, no. 3, pp. 415–437, 2002.

[23] A. W. Lo, "Efficient Markets Hypothesis," in *The New Palgrave Dictionary of Economics*, pp. 1–17, Palgrave Macmillan, London, UK, 2017.

[24] H. Dichtl and W. Drobetz, "Are stock markets really so inefficient? the case of the "halloween indicator"," *Finance Research Letters*, vol. 11, no. 2, pp. 112–121, 2013.

[25] F. AitSahlia and J.-H. Yoon, "Information stages in efficient markets," *Journal of Banking & Finance*, vol. 69, pp. 84–94, 2016.

[26] S. Kolaric and D. Schiereck, "Are stock markets efficient in the face of fear? evidence from the terrorist attacks in paris and brussels," *Finance Research Letters*, vol. 18, pp. 306–310, 2016.

[27] A. W. Lo, "Long-term Memory in Stock Market Prices," *Econometrica*, vol. 59, no. 5, pp. 1279–1313, 1991.

[28] Wikipedia, the free encyclopedia, *Efficient Markets Hypothesis*, Wikipedia, 2021.

[29] Y. Andrianto and A. R. Mirza, "A testing of efficient markets hypothesis in Indonesia stock market," *Procedia - Social and Behavioral Sciences*, vol. 219, pp. 99–103, 2016.

[30] M. R. Borges, "Efficient market hypothesis in european stock markets," *The European Journal of Finance*, vol. 16, no. 7, pp. 711–726, 2010.

[31] P. Ferreira and A. Dionisio, "How long is the memory of the us stock market?" *Physica A: Statistical Mechanics and its Applications*, vol. 451, pp. 502–506, 2016.

[32] E. Vasileiou, "Efficient markets hypothesis in the time of covid-19," *Review of Economic Analysis*, vol. 13, no. 1, pp. 45–63, 2020.

[33] M. Schatz and D. Sornette, "Inefficient bubbles and efficient drawdowns in financial markets," *International Journal of Theoretical and Applied Finance*, vol. 23, no. 7, Article ID 2050047, 2020.

[34] N. Nurdina, R. Y. Sidharta, and M. Mochklas, "Inefficient markets, anomalies, and investor behavior: a literature review," *International Journal of Economics, Business and Accounting Research (IJEBAR)*, vol. 5, no. 2, 2021.

[35] A. W. Lo, "Reconciling efficient markets with behavioral finance: the adaptive markets hypothesis," *Economic Psychology*, vol. 7, no. 2, pp. 21–44, 2004.
[36] M. Hull and F. McGroarty, "Do emerging markets become more efficient as they develop? long memory persistence in equity indices," *Emerging Markets Review*, vol. 18, pp. 45–61, 2013.
[37] J. H. Kim, A. Shamsuddin, and K.-P. Lim, "Stock return predictability and the adaptive markets hypothesis: evidence from century-long us data," *Journal of Empirical Finance*, vol. 18, no. 5, pp. 868–879, 2011.
[38] A. Urquhart and R. Hudson, "Efficient or adaptive markets? evidence from major stock markets using very long run historic data," *International Review of Financial Analysis*, vol. 28, pp. 130–142, 2013.
[39] C. M. Boya, "From efficient markets to adaptive markets: evidence from the French stock exchange," *Research in International Business and Finance*, vol. 49, pp. 156–165, 2019.
[40] P. Ferreira, "Apple, alphabet or microsoft: which is the most efficient share?" *Econometric Research in Finance*, vol. 1, no. 2, pp. 67–79, 2016.
[41] S. Ding, J. Hanna, and T. Hendershott, "How slow is the nbbo? a comparison with direct exchange feeds," *Wiley Online Library*, vol. 49, no. 2, pp. 313–332, 2014.
[42] R. P. Bartlett III and J. McCrary, "How rigged are stock markets? evidence from microsecond timestamps," *Journal of Financial Markets*, vol. 45, pp. 37–60, 2019.
[43] E. Wah, *How prevalent and profitable are latency arbitrage opportunities on us stock exchanges?*, BlackRock, Charlotte, NC, USA, 2016.
[44] M. F. Ahmed and S. Satchell, "What proportion of time is a particular market inefficient?: A method for analysing the frequency of market efficiency when equity prices follow threshold autoregressions," *Journal of Time Series Econometrics*, vol. 10, no. 2, 2016.
[45] K. R. Godfrey, "Toward a model-free measure of market efficiency," *Pacific-Basin Finance Journal*, vol. 44, pp. 97–112, 2017.
[46] Thesys Technologies, "Thesys technologies," 2020, https://www.thesystech.com/.
[47] P. Bak, C. Tang, and K. Wiesenfeld, "Self-organized criticality: an explanation of the 1/f noise," *Physical review Letters*, vol. 59, no. 4, p. 381, 1987.
[48] D. Plenz, T L. Ribeiro, S R. Miller, P A. Kells, and A. Vakili, "Self-organized criticality," *Front. Phys*, vol. 38, no. 1, p. 364, 2021.
[49] P. S. Hagan, D. Kumar, and A. Lesniewski, "Managing smile risk," *Wilmott*, vol. 1, pp. 84–108, 2002.
[50] Unlisted Trading Privileges, "Utp metrics," 2021, https://www.utpplan.com/metrics.
[51] C. Tape Association, "Cta metrics," 2021, https://www.ctaplan.com/metrics.
[52] R. A. Ariel, "High stock returns before holidays: existence and evidence on possible causes," *The Journal of Finance*, vol. 45, no. 5, pp. 1611–1626, 1990.
[53] F. J. Fabozzi, C. K. Ma, and J. E. Briley, "Holiday trading in futures markets," *The Journal of Finance*, vol. 49, no. 1, pp. 307–324, 1994.
[54] T. Chordia, R. Roll, and A. Subrahmanyam, "Market liquidity and trading activity," *The Journal of Finance*, vol. 56, no. 2, pp. 501–530, 2001.
[55] I. Tsiakas, "The economic gains of trading stocks around holidays," *Journal of Financial Research*, vol. 33, no. 1, pp. 1–26, 2010.
[56] J. Hasbrouck, "Modeling market microstructure time series," *Handbook of Statistics*, vol. 14, pp. 647–692, 1996.
[57] D. Easley and M. O'Hara, "Microstructure and asset pricing," *Handbook of the Economics of Finance*, vol. 1, pp. 1021–1051, 2003.
[58] M. O'Hara, "High frequency market microstructure," *Journal of Financial Economics*, vol. 116, no. 2, pp. 257–270, 2015.